SPORE: Combining Symmetry and Partial Order Reduction

MICHALIS KOKOLOGIANNAKIS, MPI-SWS, Germany
IASON MARMANIS, MPI-SWS, Germany
VIKTOR VAFEIADIS, MPI-SWS, Germany

Symmetry reduction (SR) and partial order reduction (POR) aim to scale up model checking by exploiting the underlying program structure: SR avoids exploring executions equivalent up to some permutation of symmetric threads, while POR avoids exploring executions equivalent up to reordering of independent instructions. While both SR and POR have been well studied individually, their combination in the context of stateless model checking has remained an open problem.

In this paper, we present SPORE, the first stateless model checker that combines SR and POR in a sound, complete and optimal manner. SPORE can leverage both symmetries in the client program itself, but also *internal symmetries* in the underlying implementation (i.e., idempotent operations), a novel symmetry notion we introduce in this paper. Our experiments confirm that SPORE explores drastically fewer executions than tools that solely employ SR/POR, thereby greatly advancing the state-of-the-art.

CCS Concepts: • **Theory of computation → Concurrency**; **Verification by model checking**.

Additional Key Words and Phrases: Model Checking, Dynamic Partial Order Reduction, Symmetry Reduction

ACM Reference Format:
Michalis Kokologiannakis, Iason Marmanis, and Viktor Vafeiadis. 2024. SPORE: Combining Symmetry and Partial Order Reduction. *Proc. ACM Program. Lang.* 8, PLDI, Article 219 (June 2024), 23 pages. https://doi.org/10.1145/3656449

1 INTRODUCTION

Stateless model checking (SMC) [Godefroid 1997] verifies a concurrent program by enumerating all of its executions. SMC is quite popular in concurrent program verification as (a) can be used by programmers without any expertise in formal methods, (b) it can handle programs in full-fledged programming languages like C, C++ and Java, and (c) it can reason about the effects of the underlying weak memory model (e.g., C/C++11 [Lahav et al. 2017]). On the downside, however, SMC only supports verification of bounded programs, and often does not scale well enough to handle client programs with a sufficient number of threads to provide strong confidence the correctness of a given implementation.

There are two sound techniques that can be employed to increase the scalability of SMC.

Symmetry reduction (SR) [Clarke et al. 1996; Emerson and Wahl 2005] exploits symmetries in the threads of the program under test (e.g., all threads running the same code) and avoids to consider all the ways in which symmetric threads interleave, as the order in which such threads execute is

Authors' addresses: Michalis Kokologiannakis, MPI-SWS, Kaiserslautern, Germany, michalis@mpi-sws.org; Iason Marmanis, MPI-SWS, Kaiserslautern, Germany, imarmanis@mpi-sws.org; Viktor Vafeiadis, MPI-SWS, Kaiserslautern, Germany, viktor@mpi-sws.org.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
© 2024 Copyright held by the owner/author(s).
ACM 2475-1421/2024/6-ART219
https://doi.org/10.1145/3656449

clearly irrelevant. As an example of SR, consider the FAIS program where N symmetric threads perform an atomic "fetch-and-increment" operation on x:

$$\texttt{fetch_add}(x, 1) \parallel \dots \parallel \texttt{fetch_add}(x, 1) \tag{FAIS}$$

While naive SMC explores $N!$ executions for this program, SR only explores 1 execution.

Dynamic partial order reduction (DPOR) [Abdulla et al. 2014; Flanagan and Godefroid 2005] reduces the program state space by not exploring executions that are equivalent up to some permutation of independent instructions (e.g., instructions accessing different variables). For instance, consider the program below where 26 (non-symmetric) threads write different parts of an array:

$$a := 1 \parallel b := 2 \parallel \dots \parallel z := 26 \tag{ARRAY}$$

For ARRAY, naive SMC would again explore 26! executions while DPOR would only explore 1, as it notices that all threads access different parts of memory, and hence their relative order is irrelevant.

A common way to view both SR and DPOR is via the *equivalence partitioning* they induce on the program state space. Indeed, SR groups together executions that can be obtained from one another by changing the ID of symmetric threads, while DPOR groups together executions that can be obtained from one another by changing the order of non-conflicting instructions.

Observe, however, that even for symmetric programs, SR and DPOR are not equivalent, and neither approach subsumes the other. This can be seen with the example below:

$$\begin{array}{l} i := \texttt{fetch_add}(x, 1) \\ a[i] := i \end{array} \parallel \dots \parallel \begin{array}{l} i := \texttt{fetch_add}(x, 1) \\ a[i] := i \end{array} \tag{FAIS+ARRAY}$$

While DPOR explores $N!$ executions for FAIS+ARRAY (due to the conflicting `fetch_add`s), SR explores $(2N-1)!!$ executions (double factorial of odd numbers). This discrepancy is because in SMC, after each thread has executed its `fetch_add`, symmetry "breaks", as each thread reads a different value.

Even though SR and DPOR are both effective when applied, existing SR/DROR approaches have two major limitations. First, they are incompatible: indeed, despite years of research on each of SR/DPOR, no algorithm manages to successfully combine the two, so employing one of them precludes the usage of the other. Second, both SR and DPOR fail to leverage *internal symmetries*, i.e., *idempotent* operations of the underlying implementation. One case of internal symmetry is the quintessential *helping* pattern, where some operation observes an ongoing operations of the same type that is incomplete, and then tries to complete the ongoing operation before performing its own. SR fails to exploit internal symmetries as the threads performing the operations are not sharing the same code, while DPOR fails to do so because the two operations are considered conflicting.

In this paper, we present SPORE (Symmetry and Partial Order Reduction Explorer), a novel algorithm that combines SR and DPOR, and overcomes both limitations above. SPORE resolves thread-level symmetries by restricting the coherence order of symmetric conflicting operations to agree with their thread order, and internal symmetries with a novel memory-model axiomatization that equates executions differing only in the order of the locally symmetric operations. The resulting algorithm is sound, complete and optimal under the combined equivalence partitionings, and achieves exponential reduction in verification time over the state-of-the-art. SPORE is also parametric in the choice of the underlying (weak) memory model.

Our contributions can be summarized as follows.

- §2 We (informally) describe why the combination of DPOR and SR is non-trivial, as well as how SPORE exploits thread-level and internal internal symmetries.
- §3 We present SPORE in detail and prove its correctness.
- §4 We implement SPORE in a tool for C/C++ programs, and empirically demonstrate that it is orders of magnitude faster than the state-of-the-art.

2 SPORE: INFORMAL DESCRIPTION

We develop Spore by adding SR on top of a DPOR algorithm (as opposed to the other way around), since DPOR underpins most modern SMC solutions [Abdulla et al. 2018; Aronis et al. 2018; Chalupa et al. 2017; Kokologiannakis et al. 2022, 2019b; Norris and Demsky 2013]. As such, we begin this section by explaining the basics of DPOR (§2.1), and then describe why the combination of DPOR and symmetry reduction is non-trivial and how Spore achieves it (§2.2). We end the section by demonstrating how Spore handles internal symmetries (§2.3).

2.1 Dynamic Partial Order Reduction

Modern DPOR algorithms, such as TruSt [Kokologiannakis et al. 2022], represent program executions up to the reordering of independent accesses in a structure called *execution graph* [Alglave et al. 2014], and verify a given program by constructing its associated execution graphs in an incremental fashion.

Each execution graph G comprises: (a) a set of events E (graph nodes), modeling instructions of the program, and (b) a few relations on events (graph edges), modeling various interactions between the instructions. In the following, we consider three such directed edges: the *program order* (po), which orders instructions of the same thread, the *reads-from* relation (rf), which relates each read event r in G to a write event w in G, from which r obtains its value, and the *coherence order* (co), which totally orders writes at each memory location.

Example 1 Consider the W+R+R program below.

$$\text{T1: } x := 1 \;\|\; \text{T2: } r_2 := x \;\|\; \text{T3: } r_3 := x \qquad \text{(W+R+R)}$$

Under sequential consistency (SC) [Lamport 1979], the program has four executions, 1–4, which model the four equivalence classes into which the $3! = 6$ thread interleavings are partitioned. These graphs can be produced by the following DPOR exploration starting from the initial graph Init through the intermediate graphs A, B, and C.

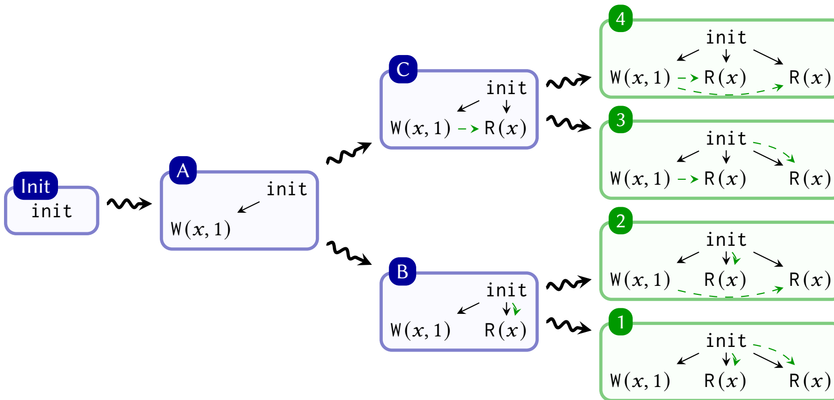

The exploration proceeds in a depth-first manner: DPOR adds the events of the program from left to right, one by one, and whenever a read has more than one place to read from, DPOR initiates a recursive subexploration. For instance, when the read of T2 is added, it can read both 0 and 1 (both options are consistent according to SC), and thus DPOR initiates subexplorations B and C. DPOR proceeds in a similar manner, until all events of the program have been added to the graph.

Conventions

Following standard conventions in the weak memory model literature, we (1) treat rf as a relation from the write to the read event; (2) assume a special initialization event init, which initializes every location with 0 and is thus po-before all other events and co-before all other write events; (3) we do not draw co edges from init to other writes (as it is trivially co-before them). In explorations, we use letters to refer to intermediate executions, numbers to refer to full executions, and red to denote executions that will not be explored.

Revisits. The exploration in Example 1 was largely straightforward, but there is still one aspect of DPOR we have not discussed: *revisiting*. For exposition purposes, suppose we add the events of W+R+R from right to left. When we encounter the reads, they cannot yet read 1 because the corresponding write does not exist in the graph. Therefore, whenever a write is added to a graph, DPOR also revisits existing same-location reads to see if they can read from the newly added write.

Whenever DPOR revisits a read r from a write w, it *restricts* the graph to remove some of the events added to the graph after r, since they may depend on the value read by r. (If not, they will be re-added in subsequent steps of the exploration.) The most common choice for restricting the graph is to keep only the events that were added before r and those causally before w (i.e., in its porf $\triangleq (\mathsf{po} \cup \mathsf{rf})^+$ prefix). For instance, in the right-to-left exploration of W+R+R, if $\mathtt{W}(x, 1)$ revisits the read of T3, the resulting graph does not have the read of T2 because it was added after T3 and is not porf-before $\mathtt{W}(x, 1)$.

The restriction due to revisits may lead to duplicate explorations, as we demonstrate below.

Example 2 Consider the following variation of W+R+R.

$$\text{T1: } x := 1 \;\big\|\; \text{T2: } r_2 := x \;\big\|\; \text{T3: } x := 2 \qquad \text{(W+R+W)}$$

Adding the events from left to right, observe that there are two subexplorations where $\mathtt{W}(x, 2)$ has the chance to revisit the read of T2: when the latter reads 0 and when it reads 1. These subexplorations are shown in Fig. 1. If $\mathtt{W}(x, 2)$ performs the revisit in both, the exact same graph will be created.

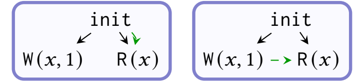

Fig. 1. Revisit opportunities

There are two ways DPOR can avoid such duplication. Abdulla et al. [2014] and Kokologiannakis et al. [2019b] simply save all encountered executions (more precisely: the ones created by revisits), and drop subsequent revisits that yield an already encountered execution. Storing executions, however, leads to exponential memory consumption in the size of the program under test.

Avoiding Duplication with Maximal Extensions. A better solution adopted by TruSt [Kokologiannakis et al. 2022] is to impose a *revisiting condition* so that a given revisit only takes place once among all possible subexplorations. The key observation is that whenever DPOR encounters two graphs that will yield the same graph immediately after a revisit, then in both cases the revisit happens from the same write w to the same read r, and the graphs only differ in the sets of events that were affected by the revisit (namely, r itself and all the events deleted by the revisit).

TruSt therefore constrains the events affected by the revisit (i.e., the read being revisited and the deleted events) to form a *maximal extension*: to be added co-maximally w.r.t. to the porf-prefix of the revisiting write. Maximal conditions are better understood with an example.

Example 3 Consider the REV-EX below along with its SC-consistent execution graphs.

$$\text{T1: } a := y \;\Big\|\; \text{T2: } \textbf{if } (x = 1)\ \ x := 2\ \ c := x \;\Big\|\; \text{T3: } x := 1\ \ y := 1 \qquad \text{(REV-EX)}$$

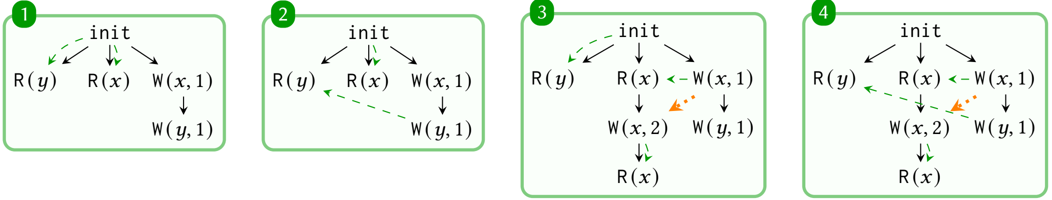

A DPOR run producing these execution can be seen below.

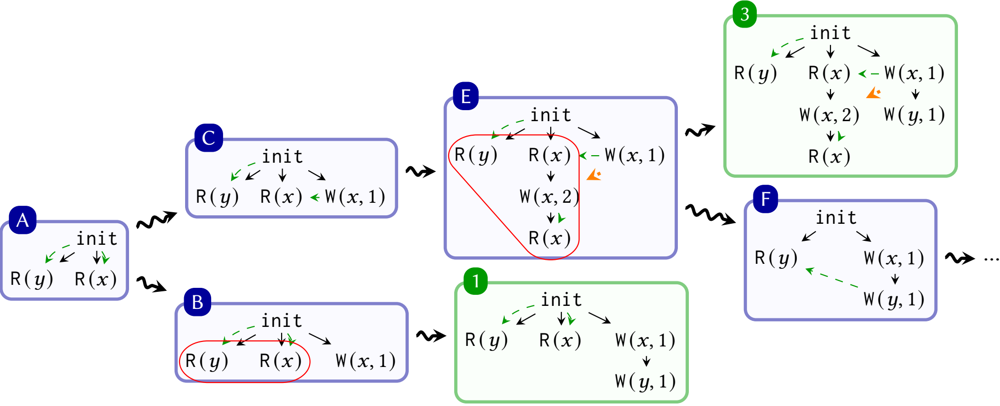

Assuming that DPOR adds events in a left-to-right manner, after adding the events of the first two threads, it then adds W$(x, 1)$ which can either revisit R(x) or not (graphs C and B, respectively).

Following the respective subexplorations, W$(y, 1)$ is encountered in both cases: in exploration B immediately, and in exploration C after adding the events under the conditional of T2[1]. Similarly to W$(x, 1)$, in both subexplorations, W$(y, 1)$ has the opportunity to either revisit R(y) or not.

Revisiting R(y) in both cases, however, leads to duplication, as the same graph (graph F) will be obtained twice. Maximal extensions dictate that the revisit only takes place from execution E, as all the affected events are added maximally w.r.t. W$(y, 1)$. To see why, it is helpful to think "backwards": starting from the graph obtained from the revisit without the write and read participating in the revisit (W$(y, 1)$ and R(y)), if all the affected events are added in a co-maximal manner (i.e., reads reading the co-latest write and writes added last in co), we get graph E, which is the graph from where the revisit takes place.

To define maximal extensions, we first introduce an auxiliary definition about execution graphs. A write event w is co-*maximal* in a set of events E if $w \in E$ and it does not have a co-successor in E (i.e., $\nexists w' \in E.\ \langle w, w' \rangle \in \text{co}$).

[1]These events have a unique co and rf option as SC enforces coherence: informally, T2 is already aware of W$(x, 1)$ so W$(x, 2)$ has to be co-after it, and R(x) has to read the latest value T2 is aware of.

Definition 2.1. An event e in a graph G is *added maximally* w.r.t. a write event w in G, if the following conditions hold, where E is the set of all events e' added before e or in w's porf prefix (i.e., $\langle e', w\rangle \in \mathsf{porf}$):

- If $e \in \mathtt{W}$, then e is co-maximal in E.
- If $e \in \mathtt{R}$, then $G.\mathsf{rf}(e)$ is co-maximal in E.

Observe that non-write/read events are always added maximally w.r.t. a revisiting write.
Maximal extensions also have the following useful property, which we will use in some of our examples below.

Proposition 2.2. *If a write w revisits a read r resulting in a graph G, the porf-prefix of w will not be removed in any of the subsequent subexplorations starting from G [Kokologiannakis et al. 2022].*

2.2 Spore: Thread-Level Symmetries

Consider again the W+R+R example where T2 and T3 share their code.

$$\text{T1: } x := 1 \;\Big|\Big|\; \text{T2: } r_2 := x \;\Big|\Big|\; \text{T3: } r_3 := x \qquad (\text{W+R+R})$$

We say that executions 2 and 3 from its consistent executions (see Example 1) are symmetric because one can be obtained by permuting the symmetric threads of the other.

2.2.1 Distinguishing Among Symmetric Executions. To avoid exploring both graphs, we pick a *representative* execution among them and instrument DPOR to drop non-representative symmetric executions.

Spore achieves this using thread IDs: we deem as representative the graph where a symmetric thread only reads values that are at least as "recent" (in terms of co) as the ones read by its symmetric predecessor. In the W+R+R example, this means that graph 2 is the representative one, as in graph 3 the read of T2 reads a value that is co-after the one read by T3[2].

Let us formalize this intuition. We say that two events e, e' in an execution graph G are *prefix-matching* (and write prefix-matching(e, e')), if they originate from threads with the same code and have matching po-prefixes, i.e., all events po-before them are either not memory accesses or reads that pairwise read from the same write. Note that two writes can be prefix-matching, but any po-later pair of events cannot be: writes break matching prefixes because they are co-ordered.

Spore picks as representative graphs the ones where the thread order of prefix-matching events does not contradict an extension of co called *extended coherence order*: $\mathsf{eco} \triangleq (\mathsf{co} \cup \mathsf{rf} \cup \mathsf{rb})^+$, where $\mathsf{rb} \triangleq \mathsf{rf}^{-1}; \mathsf{co}$ is the *reads-before* order, denoting that a read reads from a write whose value is later overwritten. Observe that, due to the definition of prefix-matching events above, any eco path between two prefix-matching events will involve co.

Given this notion of representative graphs, in the W+R+W example above, graph 2 in Example 1 is the representative because eco agrees with the thread order (there is an rb; rf path from T2 to T3), but graph 3 is not as eco contradicts the thread order.

2.2.2 Problem #1: The Interaction Between Representative and Maximal Executions. This solution, however, does not work that easily due to *revisiting* (§2.1). The problem is that SR avoids exploring certain graphs (i.e., the non-representative ones), the exploration of which DPOR might *require* so that a given revisit happens. Put differently, maximal extensions can be non-representative graphs.

[2]Recall that all writes are co-after the initializer event.

Example 4 To illustrate the problem, consider the following variation of W+R+R (again, T2 and T3 share their code), and suppose we are interested in the executions where $a = 1$.

$$\begin{array}{l|l|l|l} \text{T1: } x := 1 & \text{T2: } r_2 := x & \text{T3: } r_3 := x & \text{T4: } y := 1 \\ \quad a := y & & & \end{array} \qquad \text{(W+R+R-REV)}$$

Similarly to W+R+R, graphs 2 and 3 are symmetric, and graph 2 is the representative one.

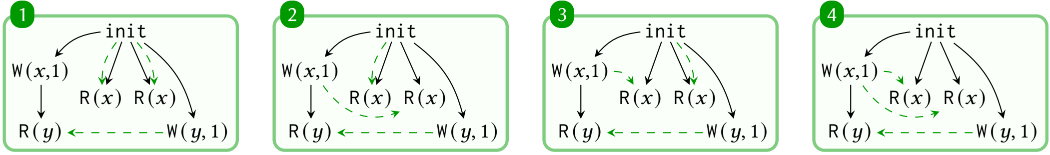

We now present a (partial) DPOR exploration of this program, with the objective of showing that the combination of DPOR and SR is not guaranteed to be correct. Concretely, we will show that execution 1 will not be generated if DPOR explores the program threads in a peculiar order[3].

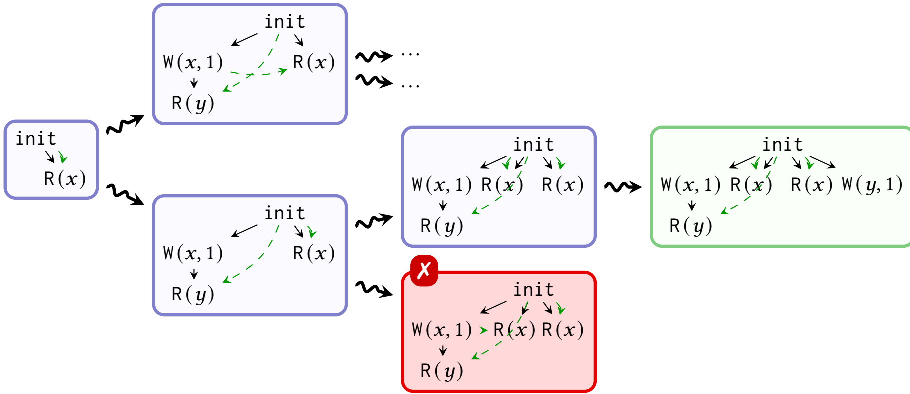

Suppose DPOR first adds the read of T3, and then proceeds with the events of T1. When it adds W$(x, 1)$, it can either revisit R(x) (top exploration tree) or not (bottom exploration tree). Since we are interested in generating execution 1, let us disregard the top exploration tree (where T3 reads 1) and focus on the bottom one. (The reason we discard the top one is that DPOR does not "undo" revisits: since W$(x, 1)$ revisits R(x) of T3, in all subsequent subexplorations T3 keep reading 1; see Prop. 2.2.)

At the next step, the algorithm will add the read of T2, which can either read 1 (from T1) or 0 (the initial value). DPOR, however, will only consider the exploration where the read is reading 0, and not the execution where it reads 1, as the latter is not the representative among the symmetric ones. (The one where T2 reads 0 and T3 reads 1 is.)

At the final step, the algorithm will add the W$(y, 1)$ event of T4, and will consider to revisit the R(y). With the maximal extension condition of §2.1, however, this revisit is doomed to fail, since the read of T2 is not added co-maximally w.r.t. W$(y, 1)$. Hence DPOR will not generate execution 1.

[3]DPOR should be able to generate all executions of a program irrespective of the order in which it encounters its threads.

As the W+R+R-REV example demonstrated, the problem when combining DPOR and SR is that resulting algorithm might deem the graphs on which TruSt's maximal extension condition enables a certain revisit as non-representative (and therefore drop them).

There are two potential solutions to this problem.

The first is to modify the maximal extension condition to hold only for representative graphs. Unfortunately, this approach does not work because of the atomicity condition of read-modify-write (RMW) operations. In our technical appendix [Kokologiannakis et al. 2024b], we show that it is impossible to define a maximality condition purely at the level of execution graphs without consulting the program.

The second solution is to *keep* the maximal extension condition intact, but restrict the *exploration order* so that representative executions always form maximal extensions. To see why restricting the exploration order is a promising solution, let us consider again Example 4. The reason why a maximal extension was created in a non-representative execution was that T3 was added before T2 (i.e., against thread order), and T2 had co-later options available to it (T1 was added after T3 but before T2). By fixing the exploration order, we essentially try to "force" co to agree with the thread order.

2.2.3 Problem #2: Fixing the Exploration Order is Inadequate. Given the above, a natural choice is to maintain a left-to-right scheduling among threads that share their code. Even though this simple modification mitigates the issue in W+R+R-REV, it does not restore correctness in general.

Example 5 To see why, consider the program below where T2 and T3 share their code, along with one of its representative executions.

T1: $a := y$	T2: $r_2 := x$	T3: $r_3 := x$	
	$x := 1$	$x := 1$	(R+RWW+RWW)
	$y := 1$	$y := 1$	

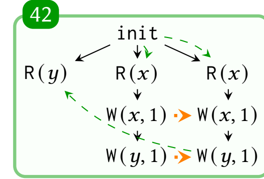

Assuming that we schedule all threads in a left-to-right manner, execution 42 cannot be generated by the procedure described so far. The first point where the algorithm has more than one choice to consider is the addition of $\mathtt{R}(x)$ of T3. The case where $\mathtt{R}(x)$ reads from $\mathtt{W}(x, 1)$ cannot lead to 42 because the restriction of the graph upon the revisit of $\mathtt{R}(y)$ will preserve the rf-edge of the $\mathtt{R}(x)$ read. Therefore, we are left with the case where $\mathtt{R}(x)$ reads from init (graph K below).

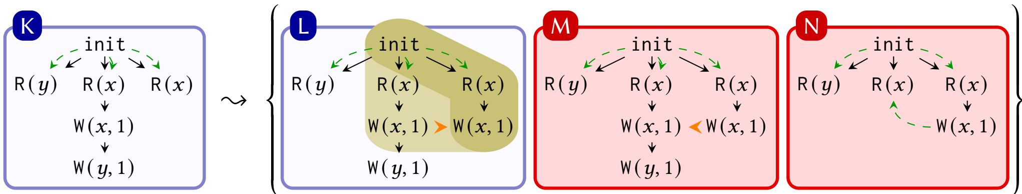

When the $\mathtt{W}(x, 1)$ of T3 is added to K, there are three options:

- L: $\mathtt{W}(x, 1)$ is added co-after T2's $\mathtt{W}(x, 1)$. This execution is explored by DPOR, but cannot lead to the graph 42 because when $\mathtt{W}(y, 1)$ is added in T3, it will be unable to revisit $\mathtt{R}(y)$ because the $\mathtt{W}(x, 1)$ of T2 is not maximally added w.r.t. T3's $\mathtt{W}(y, 1)$: it is co-before T3's $\mathtt{W}(x, 1)$, which is in T3's porf-prefix.

M: W$(x, 1)$ is co-before T2's W$(x, 1)$. This execution is dropped because co contradicts thread-order of symmetric events.

N: W$(x, 1)$ revisits the R(x) of T2. This execution is also dropped because it is not a representative one (T2 is reading a co-earlier value than T3).

As the R+RWW+RWW example above clearly demonstrates, fixing the scheduling policy is insufficient to guarantee completeness. Essentially, the issue described in §2.2.2 still persists: execution 42 could not be produced because a maximal extension was dropped (graph M) in favor of the representative one (graph L). In turn, in the representative execution L, a co-edge from a symmetric thread to the porf-prefix of the revisiting write precluded the revisit.

This last observation is key in marrying DPOR and SR: since a revisit fails due to an event of a symmetric thread being added non-maximally, Spore's solution is to consider symmetric events part of the revisiting write's prefix. In the case of R+RWW+RWW, when Spore considers the revisit between the W$(y, 1)$ of T3 and the R(y) of T1, the prefix of W$(y, 1)$ will include not just the events porf-before it, but also the porf-prefix of symmetric events as well (namely, event W$(x, 1)$ of T2). As such, graph 42 will be generated from L because all the affected events (namely, T1's R(y) and T2's W$(y, 1)$) are added maximally w.r.t. the new prefix of W$(y, 1)$.

2.2.4 Problem #3: Handling po ∪ rf ∪ co *cycles.* Changing the notion of a prefix is instrumental in restoring completeness, but comes with a caveat. In DPOR, a write can never revisit events in its own prefix. So, by introducing a new notion of a prefix (henceforth sprefix) in Spore, do we lose any executions? Is it possible that this novel notion of a prefix precludes some revisit that does not create a causal cycle, thereby rendering Spore incomplete?

The answer depends on the underlying memory model. First, we can show that sprefix cycles boil down to po ∪ rf ∪ co cycles. (Our full argument is presented in §3.) Strong models, such as SC, TSO [SPARC International Inc. 1994], and SRA [Lahav et al. 2016], forbid $(\text{po} \cup \text{rf} \cup \text{co})^{+}$ cycles, and so it is never possible for a read to read from a write in its sprefix.

In weaker models, such as RC11 [Lahav et al. 2017], however, the answer is yes: it *can* be the case that an event is in its own sprefix but not in its own porf-prefix. Such a scenario is shown below.

Example 6 Consider the SP-CYC program, where T2 and T3 share their code.

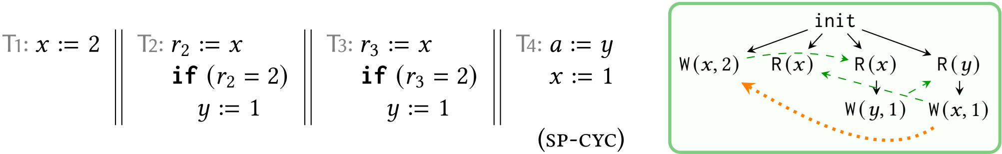

In the execution of Example 6, W$(x, 1)$ is in its own sprefix (W$(x, 1)$ is read from the R(x) of T2, which is symmetric to the R(x) of T3, which is in turn in the prefix of W$(x, 1)$), but not in its own porf-prefix (there is no porf cycle).

To restore completeness, Spore therefore checks that no consistent execution graph has a po ∪ rf ∪ co cycle. This condition typically holds: a po ∪ rf ∪ co cycle implies that there exist two writes that are not porf-ordered, and such unordered concurrent writes are rare in realistic implementations [Abdulla et al. 2019; Kokologiannakis et al. 2019b]. As we show in §4, Spore is directly applicable to realistic libraries of concurrent data structures.

```
enqueue(v) ≜
  node := malloc(...)
  node.value := v
  node.next := NULL
  do
    t := tail
    next := t.next
    if (t ≠ tail) continue
    if (next ≠ NULL)
      [CAS(tail, t, next)]
      continue
  while (¬CAS(t.next, next, node))
  [CAS(tail, t, node)]

dequeue() ≜
  do
    h := head
    n := h.next
    if (h ≠ head) continue
    if (n = NULL) return None
  while (¬CAS(head, h, n))
  t := tail
  if (h = t)
    [CAS(tail, t, n)]
  v := n.value
  reclaim(h)
  return v

rdcss_read(a₂) ≜
  r := a₂
  while (is_desc(r))
    [complete(r)]
    r := a₂
  return r

complete(d) ≜
  r := d.a₂
  n := (r = d.o₁) ?
         d.n₂ : d.o₂
  CAS(d.a₂, d, n)

rdcss(d) ≜
  r := CAS(d.a₂, d.o₂, d)
  while (is_desc(r))
    [complete(r)]
    r := CAS(d.a₂, d.o₂, d)
  if (r = d.o₂) [complete(d)]
  return r
```

Fig. 2. DGLM queue (left) and RDCSS (right). Global variables are underlined; function arguments are passed by reference; CAS returns whether it succeeded.

2.3 Spore: Internal Symmetries

We now switch gears and present how Spore exploits internal symmetries. We first present some examples of such symmetries (§ 2.3.1), and then discuss Spore's treatment (§ 2.3.2). We end this section by discussing how internal and thread-level symmetries interact (§ 2.3.3).

2.3.1 Internal Symmetry Examples. Fig. 2 shows two examples of internal symmetries: the Doherty-Groves-Luchangco-Moir (DGLM) queue [Doherty et al. 2004] and Restricted Double-Compare Single Swap (RDCSS) [Harris et al. 2002].

DGLM queue is a lock-free queue comprising two pointers *head* and *tail*. At the end of each enqueue operation, each enqueuer advances the *tail* pointer to point to the last element of the queue. If, however, a concurrent enqueuer or dequeuer detects that the *tail* pointer is lagging behind (i.e., *tail.next* ≠ NULL), it tries to advance *tail* on behalf of an incomplete enqueue.

RDCSS is a double CAS operation that takes as an argument a descriptor d containing two addresses a_1, a_2 with their expected values o_1, o_2 and a new value n_2. If both addresses contain their expected values, then the new value n_2 is stored at the second address a_2. To perform the double comparison atomically, RDCSS first tries to place its descriptor in the a_2 address, and then reads a_1 to determine whether to replace it with the new value n_2 or restore the old value o_2. In case another thread encounters the descriptor, it tries to complete the ongoing RDCSS call.

Both algorithms employ the textbook *helping pattern* [Herlihy 1991; Herlihy and Shavit 2008], where some operation A observes an ongoing, incomplete operation B and tries to complete B before performing its own. This helping pattern appears ins widely used concurrent libraries, including `libcds` [Khizhinsky n.d.], `folly` [Facebook n.d.] and `ckit` [Bahra n.d.], as well as in most algorithms described by Herlihy and Shavit [2008];

Observe that in both cases, the highlighted *main* and *helping* operations are *idempotent*: one of the CASes succeeds and all the others fail without changing the state. Moreover, their result is the same irrespective of which operation succeeds, and that the program cannot distinguish *which* operation succeeded. Indeed: (i) both operations execute exactly the same code, (ii) their returned value is not checked by the program, and (iii) swapping which of the operations succeeded preserves consistency and does not mask any error. As we will shortly see, these three conditions enable Spore to exploit internal symmetries and drastically reduce the state space. (In contrast,

thread-level symmetries are inapplicable because the main and the helping operations have different execution prefixes.)

2.3.2 Exploiting Idempotent Operations. SPORE exploits idempotent operations by only exploring executions where the *main* operation succeeds. To this end, SPORE changes the underlying memory model and treats *helping* operations as no-ops, which have no incoming/outgoing `rf` or `co` edges. To do that, SPORE requires assistance from the user: the user annotates helping operations in the program (as in Fig. 2), and then SPORE automatically treats them as no-ops and reduces the state space to be searched.

Annotations bring us to a major challenge that needs to be resolved: ensuring annotation correctness. If users incorrectly annotate a function as helping, it might mask an existing error in the user program. As such, SPORE uses a *dummy event* in the place of the function to check whether certain (sufficient) conditions hold. If they do not, SPORE reports an annotation error to the user.

Some minimal preconditions that need to hold for a function f_h to be considered as helping w.r.t. a function f_m have already been stated in §2.3.1: (i) f_h and f_m execute the same code, (ii) the returned value of f_h and f_m is not checked by the program, and (iii) replacing an execution where f_m fails and f_h succeeds with one where f_m succeeds and f_h is treated as no-op preserves consistency and the presence of an error.

Let us now go over these conditions in more detail. The first two conditions lie at the heart of idempotency, and are what allow SPORE to treat f_h as a no-op: no code uses the result of f_h and is thus safe to disregard it. Had f_h and f_m been different (or had their results been used), then annotating one of them as helping would mask errors in programs, like in the example below.

Example 7 Consider the HELPER-CF program, along with one of its execution graphs.

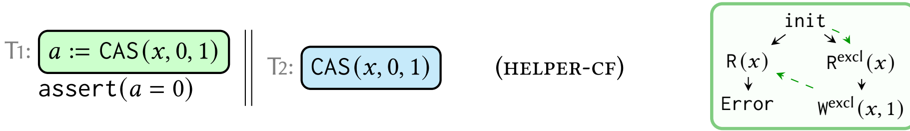

f_m and f_h are functions comprising a single CAS operation, but the result of f_m is used (i.e., f_h is incorrectly annotated as helping). If we treat f_h as a dummy event, the execution above (where the failed CAS generates a single read event and the successful one two events annotated with an excl flag) will not be explored and the error will be missed.

Condition (iii) is a bit more intricate. To ensure it, we need to guarantee that in any execution where f_h succeeds, f_m has already observed (in a synchronizing manner) the operations of f_h. If reading from writes in f_m can imply less synchronization with the rest of the program, then it is possible that reading from f_h results in an error, but reading from f_m does not (and thus, treating f_h as dummy can mask errors). We demonstrate this point with the following example.

Example 8 Consider the HELPER-SYNC program under SC.

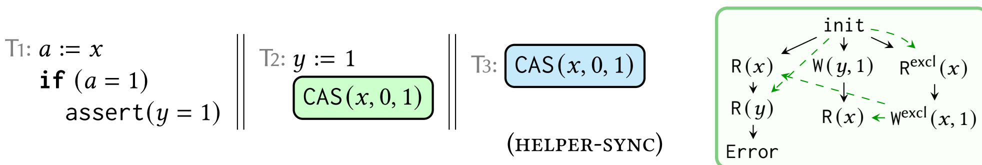

If the CAS in T2 succeeds and T1's read of x reads from it, then T1 will necessarily read $y = 1$. If, however, the CAS in T3 succeeds and T1 reads from it (as shown in the graph above), T1 can subsequently read $y = 0$ and violate its assertion (as shown in the graph above).

To fix this last issue, Spore imposes four more conditions on the user annotations:

(1) f_m and f_h have no other writes apart from a final CAS
(2) f_m has a preceding *source event* whose value it uses as the compare operand
(3) f_m is immediately preceded by a write, which is observed in a synchronizing manner before f_h
(4) all writes to the location of f_m's CAS are part of *read-modify-write* (RMW) operations

These conditions are formalized in §3. As we prove in §3, these conditions are sufficient to detect erroneously annotated helping patterns.

2.3.3 The Interaction Between Internal and Thread-Level Symmetries. Before moving on to our formal discussion of Spore, it is worth noting that idempotent operations *facilitate* SR. Consider an example with two symmetric threads performing a helping CAS. Assuming that the threads are symmetric up until the CASes, treating the CASes as an RMW operations breaks the symmetry, while treating them as dummy events preserves the symmetry.

3 SPORE: FORMAL DESCRIPTION

In this section, we describe the theoretical basis of Spore. In particular, we explain: (§ 3.1) the representation of executions as execution graphs; (§3.2) how Spore can be represented as a memory model; (§3.3) Spore's exploration algorithm; (§3.4) why Spore is correct, i.e., why it explores exactly one graph per the combined equivalence classes of DPOR and SR, and does not mask any errors.

3.1 Execution Graphs

An execution graph comprises a set of events (nodes), and a few relations on these events (edges).

Definition 3.1. An *event*, $e \in \mathsf{Event}$, is either the initialization event `init`, or a thread event $\langle t, i, l \rangle$ where $t \in \mathsf{Tid}$ is a thread identifier, $i \in \mathsf{Idx}$ is a serial number (denoting the index of an event within a thread), and $l \in \mathsf{Lab}$ is a label that takes (at least) one of the following forms:

- Write label: $\mathtt{W}^k(l, v) \in \mathtt{W}$, where k records the write attributes, $l \in \mathsf{Loc}$ the location accessed, and $v \in \mathsf{Val}$ the value written.
- Read label: $\mathtt{R}^k(l, v) \in \mathtt{R}$, where k records the read attributes, $l \in \mathsf{Loc}$ the location accessed, and $v \in \mathsf{Val}$ the value read.
- Annotated function label: $\mathtt{M}^m(f, as) \in \mathtt{M}$, where $m \in \{\mathsf{main}, \mathsf{help}\}$ is the function attribute, $f \in \mathsf{Fname}$ is the name of the function been called, and $as \in \mathsf{Val}^*$ is a sequence representing the function arguments.

Read and write attributes include the exclusivity flag excl for RMWs, and the access mode for RC11-style models. (Additional kinds of events exist for memory allocations, deallocations, assertion violations, *etc.*, but these do not affect the model checking algorithm in any meaningful way.)

Having defined events, we define execution graphs as follows.

Definition 3.2. An *execution graph* $G \in \mathsf{EXEC}$ comprises the following components:

(1) a set of events E that includes `init` and does not contain multiple events with the same thread identifier and serial number;
(2) $\mathtt{rf} : E \cap \mathtt{R} \to E \cap \mathtt{W}$, called the *reads-from* function, mapping each read event to a same-location write from where it gets its value;

(3) co $\subseteq \bigcup_{l \in \mathsf{Loc}} \mathtt{W}_l \times \mathtt{W}_l$ (where $\mathtt{W}_l \triangleq \{\mathtt{init}\} \cup \{\langle t, i, l \rangle \in E \mid l = \mathtt{W}\text{-}(l, _)\}$) called the *coherence order*, a strict partial order that is total on $\mathtt{W}_l$ for every location $l \in \mathsf{Loc}$; and
(4) $\leq$, a total order on E that represents the order in which events were incrementally added to the graph.

Conventions

We write $G.\mathsf{E}$, $G.\mathtt{rf}$, $G.\mathtt{co}$ and $\leq_G$ to project the various components of an execution graph. Given two events $e_1, e_2 \in G.\mathsf{E}$, we write $e_1 <_G e_2$ if $e_1 \leq_G e_2$ and $e_1 \neq e_2$. In relational algebra expressions, we abuse notation and write $G.\mathtt{rf}$ for the relation $\{\langle G.\mathtt{rf}(r), r \rangle \mid r \in G.\mathsf{R}\}$.
We assume that $\mathtt{init} \in \mathtt{W}$, and omit the $\emptyset$ for read/write labels with no attributes.
The functions tid, idx, loc, mod and arg respectively return the thread identifier, serial number, location, access mode and function arguments of an event, when applicable.
We write $G.\mathtt{W}$ for $G.\mathsf{E} \cap \mathtt{W}$ (and similarly for other sets), and use superscript and subscripts to restrict label sets (e.g., $\mathtt{W}_l \triangleq \{\mathtt{init}\} \cup \{w \in \mathtt{W} \mid \mathtt{loc}(w) = l\}$).

Observe that G does not have an explicit *program order* (po) component. We induce po based on our representation of events as follows:

$$\mathsf{po} \triangleq \left\{ \langle \mathtt{init}, e \rangle \;\middle|\; e \in \mathsf{Event} \setminus \{\mathtt{init}\} \right\} \cup \left\{ \langle \langle t_1, i_1, l_1 \rangle, \langle t_2, i_2, l_2 \rangle \rangle \;\middle|\; t_1 = t_2 \wedge i_1 < i_2 \right\}$$

In our technical appendix [Kokologiannakis et al. 2024b], we define two mappings from programs to sets of execution graphs: (1) $[\![.]\!]$, which ignores function annotation labels, and simply generates an event with a $\mathtt{M}^m$ label before the events corresponding to the function body; and (2) $[\![.]\!]^{\mathsf{Annot}}$, which in the case of functions annotated with help, generates only the $\mathtt{M}^{\mathsf{help}}$ event and does not generate any events for the body of the function call. Both mappings keep the rf and co components of graphs completely unconstrained. These components will be constrained by the memory model.

3.2 Consistency and Error Detection

A memory model, M, comprises three components: (a) a *causal prefix* relation, cb_{M}, (b) a *consistency predicate* $\mathsf{consistent}_{\mathsf{M}}(G)$ that determines whether an execution graph G is consistent, and (c) an $\textsc{IsErroneous}_{\mathsf{M}}(G)$ predicate, prescribing whether G contains an error (e.g., an invalid memory access) according to M.

The consistency predicate is used to constrain the semantics of a program. The annotation-ignoring (resp. annotation-aware) semantics of a program P under a memory model M, denoted $[\![P]\!]_{\mathsf{M}}$ (resp. $[\![P]\!]^{\mathsf{Annot}}_{\mathsf{M}}$), is given by the set of execution graphs in $[\![P]\!]$ (resp. $[\![P]\!]^{\mathsf{Annot}}$) that are M-consistent.

In SPORE, we assume an underlying memory model M with $\mathsf{cb}_{\mathsf{M}} = (\mathsf{po} \cup \mathtt{rf})^+$, $\mathsf{consistent}_{\mathsf{M}}(\cdot)$ being *extensible*, *prefix-closed*, and implying RMW atomicity and cb_{M}-acyclicity [Kokologiannakis et al. 2022], and $\textsc{IsErroneous}_{\mathsf{M}}(\cdot)$ being *prefix-monotone*. Models satisfying these requirements include SC [Lamport 1979], TSO [SPARC International Inc. 1994], Release-Acquire (RA) [Lahav et al. 2016], and RC11 [Lahav et al. 2017]. We then define a new memory model, SYM, with

$$\begin{aligned}
\mathsf{cb}_{\mathsf{SYM}} &\triangleq (\mathsf{po} \cup \mathtt{rf} \cup \mathsf{symb})^+ \\
\mathsf{consistent}_{\mathsf{SYM}}(G) &\triangleq \mathsf{consistent}_{\mathsf{M}}(G) \wedge \textsc{irreflexive}(\mathsf{symb}; \mathsf{eco}) \\
\textsc{IsErroneous}_{\mathsf{SYM}}(G) &\triangleq \textsc{IsErroneous}_{\mathsf{M}}(G) \vee \neg\textsc{irreflexive}((\mathsf{po} \cup \mathtt{rf} \cup \mathtt{co})^+) \\
&\quad \vee G \text{ is incorrectly annotated (see Def. 3.3 below)}
\end{aligned}$$

where $G.\mathsf{symb}$ is the *symmetry-before* order that orders prefix-matching events according to their thread order. Concretely, $\langle e_1, e_2 \rangle \in G.\mathsf{symb}$ if the following hold:

(i) $\mathtt{idx}(e_1) = \mathtt{idx}(e_2)$ and $\mathtt{tid}(e_1) < \mathtt{tid}(e_2)$
(ii) e_1 and e_2 originate from threads running the same code (and spawned consecutively),
(iii) have no preceding same-thread writes, and
(iv) for every preceding same-thread read r_1 of e_1, the corresponding (i.e., having the same index) read r_2 in $\mathtt{tid}(e_2)$ has the same $\mathtt{rf}$ (i.e., $G.\mathtt{rf}(r_1) = G.\mathtt{rf}(r_2)$).

Annotation Correctness. To ensure annotation correctness, Spore first checks that for each $f_h \in G.\mathsf{M}^{\text{HELP}}$, there exists a (unique) $f_m \in G.\mathsf{M}^{\text{main}}$ with the same arguments, and that these functions do not return any results (cf. conditions (i) and (ii) of §2.3.2), and are well-formed:.they comprise a (possibly empty) sequence of reads followed by a CAS operation, with a possible data dependency from the reads to the CAS (no other dependencies are allowed so that the locations accessed can be deduced by the arguments of f_m/f_h).

Assuming both functions has the proper form, Spore has to now ensure that (iii) holds, i.e., that their synchronization is the same. Since the definition of synchronization differs among memory models, for simplicity, we here provide a definition that works for SC and RA[4]. In what follows, we lift $\mathtt{loc}/\mathtt{exp}$ to return the location/expected-value of the CAS read following an $f_m \in G.\mathsf{M}^{\text{main}}$.

Our definition uses the notion of a *source write* s at location $\mathtt{loc}(f_m)$, which is observed before f_m (i.e., either it is po-before f_m or it is read po-before f_m), and writes the value $\mathtt{exp}(f_m)$. We also require that the immediate po-predecessor of f_m is observed before f_h, which ensures that the f_h has synchronized with everything in f_m's prefix, and that all writes to $\mathtt{loc}(f_m)$ after s are RMWs and do not write the same value as s. The latter condition ensures that f_m and f_h cannot both succeed, and that if f_h succeeds, then f_m observes its update.

Definition 3.3 (Annotation correctness). An execution G is *correctly annotated* if for all $f_h \in G.\mathsf{M}^{\text{help}}$, there exist (a) a corresponding $f_m \in G.\mathsf{M}^{\text{main}}$ with $\mathtt{arg}(f_m) = \mathtt{arg}(f_h)$ and (b) a source write $s \in G.\mathtt{W}$ with $\mathtt{loc}(s) = \mathtt{loc}(f_m)$ and $\mathtt{val}(s) = \mathtt{exp}(f_m)$ such that:

- $\langle s, f_m \rangle \in G.\mathtt{rf}^?;\mathsf{po}$ (s is observed before f_m),
- $\langle f_m, f_h \rangle \in \mathsf{po}^{-1}|_{\text{imm}};G.\mathtt{rf};\mathsf{po}$ (the immediate predecessor of f_m is observed before f_h),
- for all $w \in rng([s];\mathtt{co})$, $w \in \mathtt{W}^{\text{excl}}$ and $\mathtt{val}(w) \neq \mathtt{val}(s)$ (all subsequent writes to $\mathtt{loc}(f_m)$ are RMWs and write different values).

3.3 Exploration Algorithm

Let us now proceed by showing how Spore enumerates all SYM-consistent execution graphs of a program P. The algorithm is shown in Algorithm 1, which constructs the consistent graphs incrementally by recording the event addition order in the graphs' $\leq_G$ component. Spore is optimal in the sense that it only explores consistent execution graphs and it never explores two execution graphs that differ only in their $\leq_G$ components.

Spore verifies the input program P under a memory model M by calling Explore with the initial graph $G_\emptyset$ containing only the initialization event $\mathtt{init}$.

First, Explore(P, G) checks whether the current graph contains an error (Line 2). Note that errors are checked against Spore's memory model: they include not only errors under the underlying memory model M, but also *user annotation errors*.

In addition, recall that Spore's errors include the existence of $\mathsf{po} \cup \mathtt{rf} \cup \mathtt{co}$ cycles. Such a check is necessary to justify why exploring $\mathtt{cb}_{\mathsf{SYM}}$-acyclic execution graphs suffices: any $(\mathsf{po} \cup \mathtt{rf} \cup \mathtt{co})$-acyclic graph where the symmetry-before order does not contradict the $\mathtt{eco}$ order is also $\mathtt{cb}_{\mathsf{SYM}}$-acyclic.

[4] In our technical appendix [Kokologiannakis et al. 2024b], we provide the definition for the RC11 memory model. The definition for SC/RA is a special case of the RC11 definition.

Algorithm 1 Spore: An optimal combination of DPOR and SR

1: **procedure** $\textsc{Explore}_P(G)$
2:     **if** $\textsc{IsErroneous}_{\mathsf{SYM}}(G)$ **then exit**(“Error”)
3:     $a \leftarrow \textsc{AddNextEvent}_P(G)$
4:     **if** $a \in \mathsf{R}$ **then**
5:         **for** $w \in G.\mathsf{W}_{\mathtt{loc}(a)}$ **do** $\textsc{ExploreIfConsistent}_P(\mathsf{SetRF}(G, a, w))$
6:     **else if** $a \in \mathsf{W}$ **then**
7:         $\textsc{ExploreCOs}_P(G, a)$
8:         **for** $r \in G.\mathsf{R}_{\mathtt{loc}(a)}$ **such that** $\langle r, a\rangle \notin G.\mathsf{cb}_{\mathsf{SYM}}$ **do**
9:             $\mathit{Deleted} \leftarrow \{e \in G.\mathsf{E} \mid r <_G e <_G a \land \langle e, a\rangle \notin G.\mathsf{cb}_{\mathsf{SYM}}\}$
10:             **if** $\textsc{ShouldRevisit}(G, \langle r, a, \mathit{Deleted}\rangle)$ **then**
11:                 $\textsc{ExploreCOs}_P(\mathsf{SetRF}(G \setminus \mathit{Deleted}, r, a), a)$
12:     **else if** $a \neq \bot$ **then**
13:         $\textsc{Explore}_P(G)$

14: **procedure** $\textsc{ExploreIfConsistent}_P(G)$
15:     **if** $\mathsf{consistent}_{\mathsf{SYM}}(G)$ **then** $\textsc{Explore}_P(G)$

16: **procedure** $\textsc{ExploreCOs}_P(G, a)$
17:     **for** $w_p \in G.\mathsf{W}_{\mathtt{loc}(a)}$ **do** $\textsc{ExploreIfConsistent}_P(\mathsf{SetCO}(G, w_p, a))$

If the graph is error-free, Explore extends it by one event a from the program by calling AddNextEvent (Line 3). If there are no events to add, then a full execution of P has been explored, and Explore returns.

If a is a read, then Explore recursively explores all consistent rf options for that read. As such, for each same-location write w, Explore recursively calls itself (via the helper function ExploreIfConsistent) on the graph that results if a reads from w (Line 5). ExploreIfConsistent checks whether G is consistent (Line 15), and if so calls Explore recursively. (Recall that consistency also requires that the graph does not violate our SR principle.)

If a is a write, Spore proceeds with the non-revisit case and the revisit case, respectively. For the non-revisit case, Explore checks for all possible placements of the newly added write in co by means of ExploreCOs (Line 7).

For the revisit case, Spore also checks whether any of the existing reads of G can be *revisited* to read from a: since a was not present when their possible reads-from options were examined, Explore explores these additional rf options now. Thus, for each same-location read r that does not precede a, if revisiting r will not lead to a duplicate exploration (checked by ShouldRevisit[5]), Explore calls ExploreCOs on the graph that occurs if all the events that were added after r are deleted, excluding a and its predecessors (Line 11).

Observe, however, that as we motivated earlier in § 2.2.4, Spore only explores $\mathsf{cb}_{\mathsf{SYM}}$-acyclic execution graphs. As such, Spore never revisits reads that are in $\mathsf{cb}_{\mathsf{SYM}}$-before a (as opposed to cb_{M}-before a), as revisiting such reads would create $\mathsf{cb}_{\mathsf{SYM}}$ cycles (the $\mathsf{cb}_{\mathsf{SYM}}$-prefix of a revisiting write is always preserved).

If a has any other type (Line 13), Explore recursively calls itself.

[5] As the definition of ShouldRevisit is unnecessary for this discussion, we omit it; we refer interested readers to our technical appendix [Kokologiannakis et al. 2024b].

Remark 1. Observe that, with the exception of annotation errors, Spore does not take any special care for method annotation labels M. Indeed, this is because these are handled implicitly by the interpreter: Line 3 adds events according to our annotated semantics $[\![P]\!]^{\mathsf{Annot}}$. When the interpreter encounters a function annotated with main, it will yield an $\mathtt{M}^{\mathsf{main}}(as)$ (which is not treated specially) as well as the events of the function, while for a function annotated with help it will only yield an $\mathtt{M}^{\mathsf{help}}(as)$ event.

Remark 2. We assume that the AddNextEvent procedure (Line 3), always picks the leftmost thread among the ones that are symmetric, i.e., their next events are prefix-matching. This is necessary for the algorithm's correctness, which demands that when an event e is added, its $\mathtt{cb}_{\mathsf{SYM}}$-prefix already be present in the graph.

3.4 Soundness, Completeness and Optimality

3.4.1 Soundness of Internal Symmetries. We show that if a program P is erroneous under its standard interpretation $[\![P]\!]$ (which ignores annotations), then it is also erroneous under the annotated interpretation $[\![P]\!]^{\mathsf{Annot}}$ (which encodes annotated functions with dummy events). See [Kokologiannakis et al. 2024b] for how programs are mapped to execution graph sets.

Theorem 3.4. *Let P be an annotated program and $G \in [\![P]\!]_M$ such that IsErroneous$_M(G)$. Then, there exists $G' \in [\![P]\!]_M^{\mathsf{Annot}}$ such that IsErroneous$_{SYM}(G)$.*

Proof sketch. It suffices to show that there exists a corresponding execution G' (where every f_h being treated as a (single) dummy event $\mathtt{M}^{\mathsf{help}}(\dots)$) such that (1) IsErroneous$_{\mathsf{M}}(G')$ holds, or (2) G' is incorrectly annotated (see Def. 3.3). The lack of an annotation error is essential in showing that changing G' so that f_m succeeds instead of f_h does not affect G's consistency.

The conditions of Def. 3.3 essentially enforce that in any execution where f_h would succeed, (a) there is an f_m, running the same code, (b) f_m fails (there can only be one write that writes the expected value), (c) f_m reads from the CAS of f_h, or from a co-later (due to coherence and the presence of the source event), and therefore there is a porf-path from the CAS of f_h to the CAS of f_m (all writes to the CAS location are part of an RMW, and thus such a co path is also a porf path), (d) f_m is preceded by a write that was observed by the thread of f_h. This guarantees that swapping the events of f_m with those of f_h, and replacing the events of f_h with a no-op, adds no synchronization in the execution, and therefore preserves both consistency and the presence of an error.

If any of the previous conditions fails, we show that there exists an execution with f_h being treated as a no-op that is not correctly annotated. □

3.4.2 Correctness of Spore. To state our desired result, we first need to formally define which are the execution graphs that are considered equivalent up to symmetry. Given a program P with N threads, a *valid thread permutation* π is a bijection $\{1, \dots, N\} \mapsto \{1, \dots, N\}$ such that threads $\pi(i)$ and i share the same code for all $1 \leq i \leq N$. We say that two executions G_1 and G_2 are *symmetric*, denoted $G_1 \approx G_2$, if there exists a valid thread permutation π such that $\pi(G_1) = G_2$, where $\pi(G_1)$ applies the permutation to all the thread IDs in the events of G_1.

The following proposition demonstrates that the class of M-consistent execution graphs up to symmetry corresponds (one-to-one) to the class of SYM-consistent execution graphs.

Proposition 3.5. *Given a program P and an execution graph $G \in [\![P]\!]_M^{\mathsf{Annot}}$, there is a unique execution graph $G' \in [\![P]\!]_{SYM}^{\mathsf{Annot}}$ such that $G \approx G'$.*

Proof. To obtain G' from G, sort the threads running the same function by the eco of the respective events (lexicographically, in po order). It is easy to see that this ordering is well-defined

(there are no cycles), and unique: any possibly eco-unordered threads are in fact equal, and that the constructed graph G' satisfies $\textsc{irreflexive}(\texttt{symb}; \texttt{eco})$. □

Correctness of the exploration algorithm follows by adapting the proof of Awamoche [Kokologiannakis et al. 2023] and is captured by the following proposition.

Proposition 3.6 (Algorithmic Correctness and Optimality).

(1) $\textsc{Explore}_P(G_\emptyset)$ *terminates.*
(2) $\textsc{Explore}_P(G_\emptyset)$ *only explores* $\texttt{cb}_{SYM}$*-prefixes of executions in* $[\![P]\!]^{\text{Annot}}_{SYM}$.
(3) $\textsc{Explore}_P(G_\emptyset)$ *explores every execution* $G \in [\![P]\!]^{\text{Annot}}_{SYM}$ *such that* $\textsc{irreflexive}(G.\texttt{cb}_{SYM})$.
(4) $\textsc{Explore}_P(G_\emptyset)$ *never explores the same* G *twice.*

Termination holds because either a revisit step is performed and the part of the graph that cannot be changed grows or a non-revisit step is performed and the execution graph grows. Soundness holds by construction because consistency is checked before every recursive call. Completeness is more elaborate: it holds because all possible rf/co options are considered for each newly added event, and moreover previous reads can be revisited in their maximal extension (which always exists and is consistent). Optimality holds because there cannot be two steps leading to the same graph; in case of revisits, that is precluded by the uniqueness of maximal extensions.

We next show that if $[\![P]\!]^{\text{Annot}}_{\text{SYM}}$ includes a $\texttt{cb}_{\text{SYM}}$-cyclic execution, which the algorithm would not explore, then it also includes a $\texttt{cb}_{\text{SYM}}$-acyclic execution with a $\texttt{po} \cup \texttt{rf} \cup \texttt{co}$ cycle, which the algorithm would explore and report.

Proposition 3.7 ($\texttt{cb}_{\text{SYM}}$ cycle). *If there is an execution* $G \in [\![P]\!]^{\text{Annot}}_{SYM}$ *with a* $G.\texttt{cb}_{SYM}$ *cycle, then there is an execution* $G' \in [\![P]\!]^{\text{Annot}}_{SYM}$ *such that* $\textsc{irreflexive}(G'.\texttt{cb}_{SYM})$ *and* G' *has a* $\texttt{po} \cup \texttt{rf} \cup \texttt{co}$ *cycle.*

Combining Prop. 3.5, Prop. 3.6(3), and Prop. 3.7, we obtain our completeness result.

Theorem 3.8 (Completeness). *If there exists* $G \in [\![P]\!]^{\text{Annot}}_{SYM}$ *such that* $\textsc{IsErroneous}_{SYM}(G)$*, then* $\textsc{Explore}_P(G_\emptyset)$ *will report an error. Otherwise, for each* $G \in [\![P]\!]^{\text{Annot}}_{M}$, $\textsc{Explore}_P(G_\emptyset)$ *will explore an execution* $G' \in [\![P]\!]^{\text{Annot}}_{SYM}$ *such that* $G \approx G'$.

Combining Prop. 3.5 and Prop. 3.6(4), we obtain our optimality result.

Theorem 3.9 (Optimality). *For any two executions* G *and* G' *explored by* $\textsc{Explore}_P(G_\emptyset)$, $G \not\approx G'$.

4 EVALUATION

We implemented Spore as a tool for C/C++ programs on top of the open-source GenMC stateless model checker, which implements the TruSt algorithm for DPOR. We reused GenMC's infrastructure for interpreting programs and constructing and maintaining execution graphs, but replaced GenMC's consistency checking and error detection mechanism with the ones described in §3.1. We also modified the notion of a prefix used in graph construction to use $\texttt{cb}_{\text{SYM}}$, and made GenMC's scheduler respect $\texttt{cb}_{\text{SYM}}$ when encountering symmetric threads.

4.1 Goals

We evaluate Spore on a set of real-world implementations with two goals: (1) show that Spore scales well enough to verify useful implementations (and determine its scalability limit), and (2) determine to what extent its scalability should be attributed to internal vs thread-level symmetries.

To attain these goals, we run Spore on a set of representative real-world clients and benchmarks. The clients evaluate the effectiveness of the SR algorithm, while the benchmarks evaluate the effectiveness of Spore's modeling of internal symmetries. To further study how internal and

thread-level symmetries contribute to Spore's performance, we compares Spore against (a) plain SMC enhanced with SR (SR), (b) a baseline TruSt implementation (TruSt), (c) Spore without thread-level symmetries (DPOR+IS), and (d) Spore without internal symmetries (DPOR+SR). Our evaluation is performed under RC11.

As we show, Spore yields a huge improvement over the state-of-the-art as it can gracefully scale to up to 6 threads (often to many more), and both internal and thread-level symmetries are crucial for its scalability to more threads.

Experimental Setup. We conducted all experiments on a Dell PowerEdge R6525 system running a custom Debian-based distribution with 2 AMD EPYC 7702 CPUs (256 cores @ 2.80 GHz) and 2TB of RAM. We set the timeout limit to 30 minutes (denoted by ⌚). All times are in seconds.

We also ran some of our benchmarks against the DPOR implementation of Nidhugg [Abdulla et al. 2014], which obtained similar and/or worse results than TruSt (see [Kokologiannakis et al. 2024b]).

4.2 Benchmarks

To evaluate the effectiveness of thread-level symmetries, we used three different clients:

- Multiset(N): $\lceil \frac{N}{2} \rceil$ (resp. $\lfloor \frac{N}{2} \rfloor$) threads insert (resp. remove) elements at a data structure; the client checks whether each removed element was previously inserted.
- LIFO/FIFO(N): two threads check for the LIFO/FIFO property, while $\lceil \frac{N}{2} \rceil$ (resp. $\lfloor \frac{N}{2} \rfloor$) threads create "noise" in the queue to increase traffic, by inserting (resp. removing) elements.
- Empty(N): N threads insert an element and subsequently remove an element; the client ensures each removal succeeds.

As it can be seen, the clients become progressively more challenging in the sense that the number of multiple operations per thread increases, which hinders symmetry reduction.

To demonstrate that Spore is applicable to non-data-structure benchmarks as well, we used two other clients (Fig. 4):

- Mutex(N): N threads perform a lock followed by an unlock operation.
- RDCSS(N): N threads perform an RDCSS call followed by an RDCS/read call, and 2 threads perform a single RDCSS call.

To evaluate the effectiveness of internal symmetries, we used some representative benchmarks both with and without idempotent operations:

- `msqueue` [Michael and Scott 1998], `dglmqueue` [Doherty et al. 2004], `folqueue` [Fober et al. 2001] and `rdcss` [Harris et al. 2002] all employ idempotent operations.
- `treiber` [Treiber 1986], `ttaslock` [Herlihy and Shavit 2008, §7.2] and `twalock` [Dice and Kogan 2019] do not employ idempotent operations.

These benchmarks exercise different aspects of internal symmetries so that the individual effects of each symmetry type are more visible.

We also note that we have identified idempotent operations in various widely used concurrency libraries (e.g., `libcds` [Khizhinsky n.d.], `folly` [Facebook n.d.], `ckit` [Bahra n.d.]). Even though Spore's support for C++ precluded us from using `libcds` and `folly` as benchmarks, we did manage to run certain benchmarks from `ckit`, with similar performance gains.

4.3 Results

Our results are summarized in Fig. 3[6]. First, as explained in §1, SR alone is inadequate for scalability, and using a combination of DPOR and SR is crucial: with the exception of a few benchmarks, SR

[6]Detailed tables can be found in [Kokologiannakis et al. 2024b].

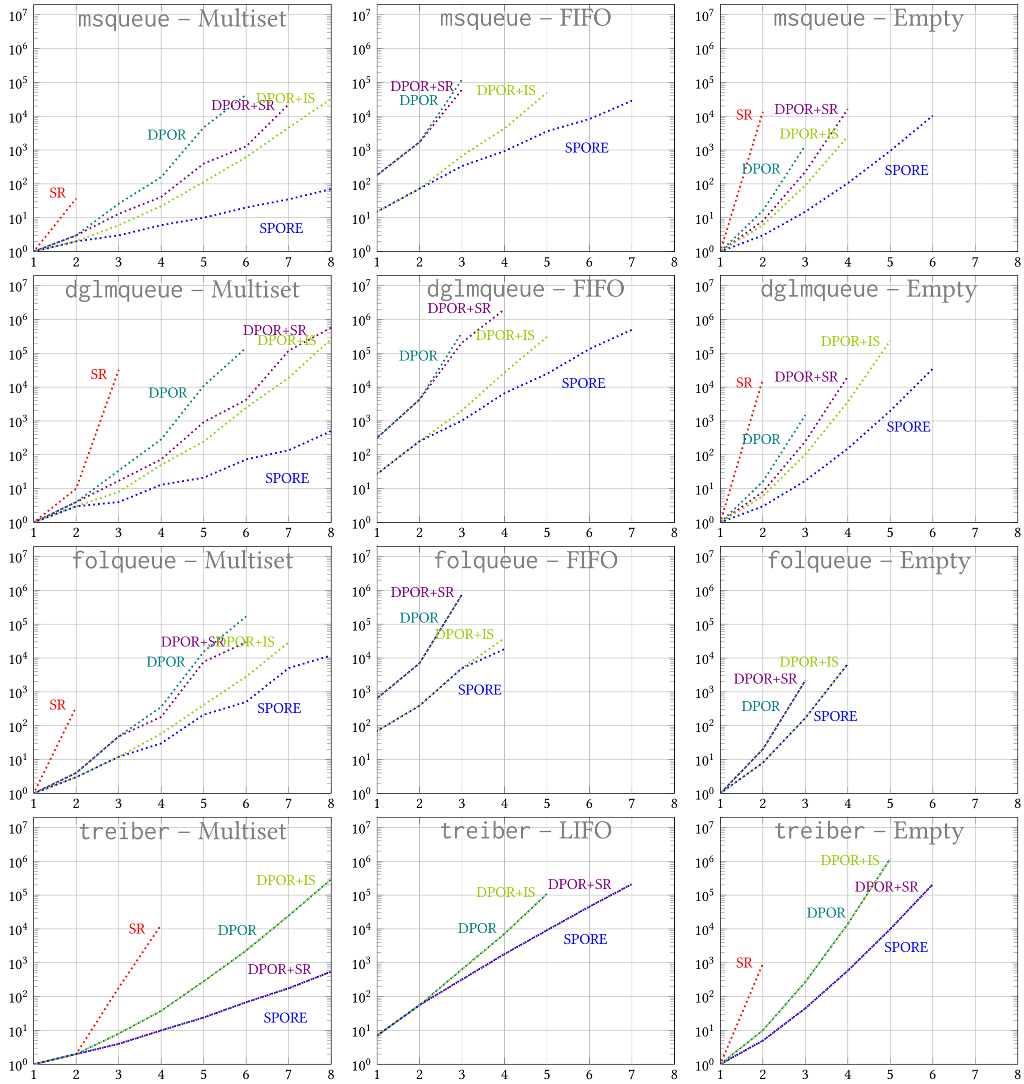

Fig. 3. Data structure benchmarks: Number of executions expored (Y-axis) per input parameter (X-axis)

consistently times out (and we therefore dismiss it for the rest of this discussion). Second, both thread-level and internal symmetries are crucial for scaling to more threads: exclusively either kind of symmetry typically leads to timeouts for some number of threads.

Let us now examine the benchmarks in more detail, starting with the multiset client (left column). The main takeaway from this client is immediately evident: while TruSt typically scales up to 6 threads before timing out, Spore scales gracefully to 8 threads (and more). Looking more closely, however, there are a few other interesting aspects as well.

Starting with `msqueue` and `dglmqueue`[7], TruSt times out for 6 threads and above, while Spore can scale up to many more. The reason for that is simple: the CAS instruction present in the queue's

[7]These benchmarks only differ in their dequeue method, which is why the results are very similar.

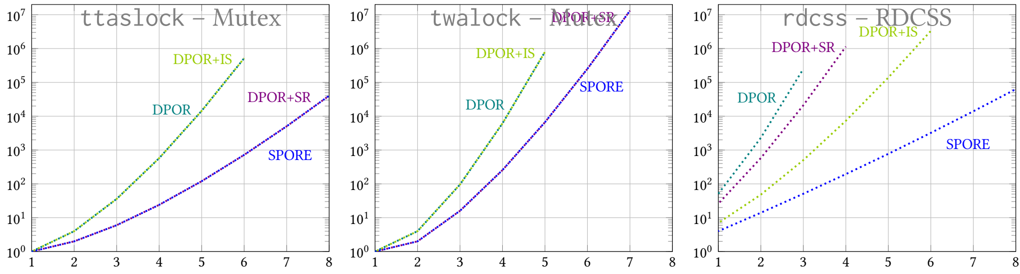

Fig. 4. Non-data-structure benchmarks: Number of executions expored (Y-axis) per input parameter (X-axis)

idempotent operation breaks symmetry, thereby leading to state-space explosion. Spore, on the other hand, runs lickety-split: it explores a single execution when the client is fully symmetric (up to 4 threads), and a small number of executions otherwise (modeling the different ways insertions interfere with deletions). As the number of dequeuers increases, Spore explores more executions,as there are more ways for deletions to interfere with insertions.

Moving on to `folqueue` and `treiber`, we can make observations similar to the ones for the previous benchmarks, albeit a bit toned down. In the case of `folqueue`, thread-level symmetries have a limited effect, as each thread uses a distinct (global) location to dispose pointers, which breaks symmetry among threads early: Spore performs similarly to DPOR+IS, while TruSt performs similarly to DPOR+SR. Analogously, in `treiber`, internal symmetries have no effect, as the code has no idempotent operations: Spore performs just as well as DPOR+SR, while DPOR+IS performs just as well as TruSt.

Generally, we observe that DPOR+IS performs better than DPOR+SR in the multiset client when both thread-level and internal symmetries are present, implying that internal symmetries carry more weight when it comes to scaling to more threads. This should not come as a surprise. Idempotent operations might be performed more than once per thread, while thread-level symmetry will break after the first non-symmetric operation. As such, since the number of idempotent operations is greater than the number of threads, internal symmetries offer a greater reduction.

Next, we move on to the other two clients. In a similar fashion, Spore scales much better than TruSt (which only manages to terminate within the time limit for two or three configurations), although it does not manage to finish within the time limit for all configurations, since these clients are not completely symmetric (like the multiset one). As expected, Spore performs better in the LIFO/FIFO (where it can better leverage the symmetry in the client), and DPOR+IS performs better than DPOR+SR whenever there are internal symmetries, for the same reasons as in the multiset client. (Note that Spore performs similarly to DPOR+IS for the first configuration of each benchmark in the LIFO/FIFO client, as SR requires at least two symmetric threads to have any effect.)

Finally, in Fig. 4 we compare all tools on some non-data-structure benchmarks. The two locking benchmarks do not employ idempotent operations, and thus Spore coincides with DPOR+SR, which has an exponentially smaller state-space than plain DPOR. In contrast, `rdcss` makes heavy use of idempotent operations, and so Spore manages to scale way better than plain DPOR.

5 RELATED WORK

As far as symmetry reduction is concerned, it has mostly been explored in the context of stateful model checking [Clarke et al. 1996; Emerson and Wahl 2005; Wahl and Donaldson 2010]. In that setting, the main challenge is to identify when two threads are symmetric, that is computationally

as hard as the graph isomorphism problem. By contrast, Spore is able to detect when two threads are symmetric on-the-fly, though in principle the reductions it achieves are not as good as the ones in stateful model checking.

As far as internal symmetries are concerned, even though a lot of effort has been devoted into making DPOR algorithms more efficient and scalable during the past few years (e.g., [Abdulla et al. 2015, 2017, 2018; Aronis et al. 2018; Chalupa et al. 2017; Chatterjee et al. 2019; Kokologiannakis et al. 2017, 2022, 2019b; Nguyen et al. 2018; Norris and Demsky 2013; Rodríguez et al. 2015]), most works focus on improving the core of DPOR and do not take into consideration the programs under test. SAVer [Kokologiannakis et al. 2021] and LAPOR [Kokologiannakis et al. 2019a] extend DPOR for programs that have spinloops and locks, respectively, while constrained-DPOR [Albert et al. 2018] takes programmer annotations into account in order to consider certain atomic operations non-conflicting.

In a different context, there has been a large body of work on static verification of concurrent programs, with techniques such as bounded model checking (BMC) or abstraction-based techniques (e.g., [Clarke et al. 2004; Elmas et al. 2009; Flanagan et al. 2005; Gavrilenko et al. 2019]). We expect that—at least for SAT/SMT-based techniques—both thread-level and internal symmetries could be exploited in a similar fashion to reduce the size of the resulting SAT formula and speed up the verification.

6 CONCLUSION

We presented Spore, a novel model checking algorithm that combines DPOR with symmetry reduction, and also exploits internal symmetries of C/C++ concurrent data structures. Our experiments confirm that Spore outperforms the state-of-the-art by a wide margin.

There are several ways this work could be extended. First, we would like to see whether Spore can handle other classes of programs in related domains, namely distributed algorithms and/or persistent programs, where similar symmetries appear. It remains to be seem whether those patterns exhibit symmetries that can be exploited in a similar fashion to enhance the applicability of automated verification techniques in those domains. Second, it would also be interesting whether Spore can be applied to models like ARMv8 [Flur et al. 2016] and POWER [Alglave et al. 2014] that do allow TruSt's $\mathsf{po} \cup \mathtt{rf}$ cycles in consistent executions (which Spore does not currently produce). Finally, Spore could also be combined with testing techniques, so that only representative executions are produced when obtaining traces of a concurrent program.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable feedback. This work was supported by a European Research Council (ERC) Consolidator Grant for the project "PERSIST" under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101003349).

DATA-AVAILABILITY STATEMENT

The benchmarks and tools used to produce the results of this paper can be found at [Kokologiannakis et al. 2024a]. Spore is available at [Kokologiannakis n.d.].

REFERENCES

Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas. 2015. "Stateless model checking for TSO and PSO." In: *TACAS 2015* (LNCS). Vol. 9035. Springer, Berlin, Heidelberg, 353–367. https://doi.org/10.1007/978-3-662-46681-0_28.

Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. 2014. "Optimal dynamic partial order reduction." In: *POPL 2014*. ACM, New York, NY, USA, 373–384. https://doi.org/10.1145/2535838.2535845.

Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. Sept. 2017. "Source sets: A foundation for optimal dynamic partial order reduction." *J. ACM*, 64, 4, (Sept. 2017), 25:1–25:49. https://doi.org/10.1145/3073408.

Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, Magnus Lång, Tuan Phong Ngo, and Konstantinos Sagonas. Oct. 10, 2019. "Optimal stateless model checking for reads-from equivalence under sequential consistency." *Proc. ACM Program. Lang.*, 3, (Oct. 10, 2019), 150:1–150:29, OOPSLA, (Oct. 10, 2019). https://doi.org/10.1145/3360576.

Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Tuan Phong Ngo. Oct. 2018. "Optimal stateless model checking under the release-acquire semantics." *Proc. ACM Program. Lang.*, 2, OOPSLA, (Oct. 2018), 135:1–135:29. https://doi.org/10.1145/3276505.

Elvira Albert, Miguel Gómez-Zamalloa, Miguel Isabel, and Albert Rubio. 2018. "Constrained dynamic partial order reduction." In: *CAV 2018*. Ed. by Hana Chockler and Georg Weissenbacher. Springer International Publishing, Cham, 392–410. ISBN: 978-3-319-96142-2. https://doi.org/10.1007/978-3-319-96142-2_24.

Jade Alglave, Luc Maranget, and Michael Tautschnig. July 2014. "Herding cats: Modelling, simulation, testing, and data mining for weak memory." *ACM Trans. Program. Lang. Syst.*, 36, 2, (July 2014), 7:1–7:74. https://doi.org/10.1145/2627752.

Stavros Aronis, Bengt Jonsson, Magnus Lång, and Konstantinos Sagonas. 2018. "Optimal dynamic partial order reduction with observers." In: *TACAS 2018* (LNCS). Vol. 10806. Springer, 229–248. https://doi.org/10.1007/978-3-319-89963-3_14.

Samy Al Bahra. N.d. *Concurrency Kit.* (). https://github.com/concurrencykit/ck.

Marek Chalupa, Krishnendu Chatterjee, Andreas Pavlogiannis, Nishant Sinha, and Kapil Vaidya. Dec. 2017. "Data-centric dynamic partial order reduction." *Proc. ACM Program. Lang.*, 2, POPL, (Dec. 2017), 31:1–31:30. https://doi.org/10.1145/3158119.

Krishnendu Chatterjee, Andreas Pavlogiannis, and Viktor Toman. Oct. 2019. "Value-Centric Dynamic Partial Order Reduction." *Proc. ACM Program. Lang.*, 3, OOPSLA, (Oct. 2019). https://doi.org/10.1145/3360550.

Edmund M. Clarke, Somesh Jha, Reinhard Enders, and Thomas Filkorn. 1996. "Exploiting symmetry in temporal logic model checking." *Form. Meth. Syst. Des.*, 9, 1/2, 77–104. https://doi.org/10.1007/BF00625969.

Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. 2004. "A tool for checking ANSI-C programs." In: *TACAS 2004* (LNCS). Vol. 2988. Springer, Berlin, Heidelberg, 168–176. https://doi.org/10.1007/978-3-540-24730-2_15.

Dave Dice and Alex Kogan. 2019. "TWA – Ticket Locks Augmented with a Waiting Array." In: *Euro-Par 2019*. Springer-Verlag, Berlin, Heidelberg, 334–345. ISBN: 978-3-030-29399-4. https://doi.org/10.1007/978-3-030-29400-7_24.

Simon Doherty, Lindsay Groves, Victor Luchangco, and Mark Moir. 2004. "Formal Verification of a Practical Lock-Free Queue Algorithm." In: *FORTE 2004* (LNCS). Ed. by David de Frutos-Escrig and Manuel Núñez. Vol. 3235. Springer, 97–114. https://doi.org/10.1007/978-3-540-30232-2_7.

Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. 2009. "A calculus of atomic actions." In: *POPL 2009*. Ed. by Zhong Shao and Benjamin C. Pierce. ACM, 2–15. https://doi.org/10.1145/1480881.1480885.

E. Allen Emerson and Thomas Wahl. 2005. "Dynamic Symmetry Reduction." In: *TACAS 2005* (LNCS). Ed. by Nicolas Halbwachs and Lenore D. Zuck. Vol. 3440. Springer, 382–396. https://doi.org/10.1007/978-3-540-31980-1_25.

Facebook. N.d. *Folly: Facebook Open-source Library.* (). https://github.com/facebook/folly.

Cormac Flanagan, Stephen N. Freund, and Shaz Qadeer. 2005. "Exploiting Purity for Atomicity." *IEEE Trans. Software Eng.*, 31, 4, 275–291. https://doi.org/10.1109/TSE.2005.47.

Cormac Flanagan and Patrice Godefroid. 2005. "Dynamic partial-order reduction for model checking software." In: *POPL 2005*. ACM, New York, NY, USA, 110–121. https://doi.org/10.1145/1040305.1040315.

Shaked Flur, Kathryn E. Gray, Christopher Pulte, Susmit Sarkar, Ali Sezgin, Luc Maranget, Will Deacon, and Peter Sewell. 2016. "Modelling the ARMv8 architecture, operationally: Concurrency and ISA." In: *POPL 2016*. ACM, St. Petersburg, FL, USA, 608–621. ISBN: 978-1-4503-3549-2. https://doi.org/10.1145/2837614.2837615.

Dominique Fober, Yann Orlarey, and Stéphane Letz. 2001. *Optimised Lock-Free FIFO Queue.* Technical Report. GRAME. https://hal.archives-ouvertes.fr/hal-02158792.

Natalia Gavrilenko, Hernán Ponce-de-León, Florian Furbach, Keijo Heljanko, and Roland Meyer. 2019. "BMC for weak memory models: Relation analysis for compact SMT encodings." In: *CAV 2019*. Ed. by Isil Dillig and Serdar Tasiran. Springer International Publishing, Cham, 355–365. ISBN: 978-3-030-25540-4. https://doi.org/10.1007/978-3-030-25540-4_19.

Michalis Kokologiannakis. N.d. *GenMC: Generic model checking for C programs.* (). https://github.com/MPI-SWS/genmc.

Patrice Godefroid. 1997. "Model checking for programming languages using VeriSoft." In: *POPL 1997*. ACM, Paris, France, 174–186. https://doi.org/10.1145/263699.263717.

Timothy L. Harris, Keir Fraser, and Ian A. Pratt. 2002. "A Practical Multi-word Compare-and-Swap Operation." In: *DISC 2002* (LNCS). Ed. by Dahlia Malkhi. Vol. 2508. Springer, 265–279. https://doi.org/10.1007/3-540-36108-1_18.

Maurice Herlihy. 1991. "Wait-Free Synchronization." *ACM Trans. Program. Lang. Syst.*, 13, 1, 124–149.

Maurice Herlihy and Nir Shavit. 2008. *The art of multiprocessor programming.*

Max Khizhinsky. N.d. *CDS C++ library.* (). https://github.com/khizmax/libcds.

Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. Dec. 2017. “Effective stateless model checking for C/C++ concurrency.” *Proc. ACM Program. Lang.*, 2, POPL, (Dec. 2017), 17:1–17:32. https://doi.org/10.1145/3158105.

Michalis Kokologiannakis, Iason Marmanis, Vladimir Gladstein, and Viktor Vafeiadis. Jan. 2022. “Truly stateless, optimal dynamic partial order reduction.” *Proc. ACM Program. Lang.*, 6, POPL, (Jan. 2022). https://doi.org/10.1145/3498711.

Michalis Kokologiannakis, Iason Marmanis, and Viktor Vafeiadis. June 2024a. *SPORE: Combining Symmetry and Partial Order Reduction (Replication Package).* (June 2024). https://doi.org/10.5281/zenodo.10798179.

Michalis Kokologiannakis, Iason Marmanis, and Viktor Vafeiadis. June 2024b. “Spore: Combining Symmetry and Partial Order Reduction (supplementary material),” (June 2024). https://plv.mpi-sws.org/genmc.

Michalis Kokologiannakis, Iason Marmanis, and Viktor Vafeiadis. 2023. “Unblocking Dynamic Partial Order Reduction.” In: *CAV 2023.* Vol. 13964. Springer, 230–250. https://doi.org/10.1007/978-3-031-37706-8_12.

Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. Oct. 2019a. “Effective lock handling in stateless model checking.” *Proc. ACM Program. Lang.*, 3, OOPSLA, (Oct. 2019). https://doi.org/10.1145/3360599.

Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019b. “Model checking for weakly consistent libraries.” In: *PLDI 2019.* ACM, New York, NY, USA. https://doi.org/10.1145/3314221.3314609.

Michalis Kokologiannakis, Xiaowei Ren, and Viktor Vafeiadis. 2021. “Dynamic Partial Order Reductions for Spinloops.” In: *FMCAD 2021.* IEEE, 163–172. https://doi.org/10.34727/2021/isbn.978-3-85448-046-4_25.

Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. 2016. “Taming Release-acquire Consistency.” In: *POPL 2016.* ACM, St. Petersburg, FL, USA, 649–662. ISBN: 978-1-4503-3549-2. https://doi.org/10.1145/2837614.2837643.

Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. “Repairing sequential consistency in C/C++11.” In: *PLDI 2017.* ACM, Barcelona, Spain, 618–632. ISBN: 978-1-4503-4988-8. https://doi.org/10.1145/3062341.3062352.

Leslie Lamport. Sept. 1979. “How to Make a Multiprocessor Computer that Correctly Executes Multiprocess Programs.” *IEEE Trans. Computers*, 28, 9, (Sept. 1979), 690–691. https://doi.org/10.1109/TC.1979.1675439.

Maged M. Michael and Michael L. Scott. 1998. “Nonblocking algorithms and preemption-safe locking on multiprogrammed shared memory multiprocessors.” *J. Parallel Distrib. Comput.*, 51, 1, 1–26.

Huyen T. T. Nguyen, César Rodríguez, Marcelo Sousa, Camille Coti, and Laure Petrucci. 2018. “Quasi-optimal partial order reduction.” In: *CAV 2018* (LNCS). Ed. by Hana Chockler and Georg Weissenbacher. Vol. 10982. Springer, 354–371. https://doi.org/10.1007/978-3-319-96142-2_22.

Brian Norris and Brian Demsky. 2013. “CDSChecker: Checking concurrent data structures written with C/C++ atomics.” In: *OOPSLA 2013.* ACM, 131–150. https://doi.org/10.1145/2509136.2509514.

César Rodríguez, Marcelo Sousa, Subodh Sharma, and Daniel Kroening. 2015. “Unfolding-based Partial Order Reduction.” In: *CONCUR 2015* (LIPIcs). Vol. 42. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 456–469. https://doi.org/10.4230/LIPIcs.CONCUR.2015.456.

SPARC International Inc.. 1994. *The SPARC architecture manual (version 9).* Prentice-Hall.

R. Kent Treiber. 1986. *Systems Programming: Coping with Parallelism.* Tech. rep. Technical Report RJ5118, IBM. https://domino web.draco.res.ibm.com/58319a2ed2b1078985257003004617ef.html.

Thomas Wahl and Alastair Donaldson. 2010. “Replication and Abstraction: Symmetry in Automated Formal Verification.” 2, 2, 799–847. https://doi.org/10.3390/sym2020799.

Received 2023-11-16; accepted 2024-03-31